Balancing performance and complexity with adaptive graph coarsening

Marek Dědič[1 2] Lukáš Bajer[2] Pavel Procházka[2] Martin Holeňa[3]

[1]Faculty of Nuclear Sciences and Physical Engineering, Czech Technical University in Prague

[2]Cisco Systems, Inc. [3]Institute of Computer Science, Czech Academy of Sciences

Motivation

Graph neural networks (GNNs) have proven to achieve superior performance on a number of graph datasets and are adopted across many fields. Superior GNN performance is, however, often paid for by high computational demands, which may be a prohibitive issue in some applications. We investigate the interplay of graph coarsening and the performance of a downstream task.

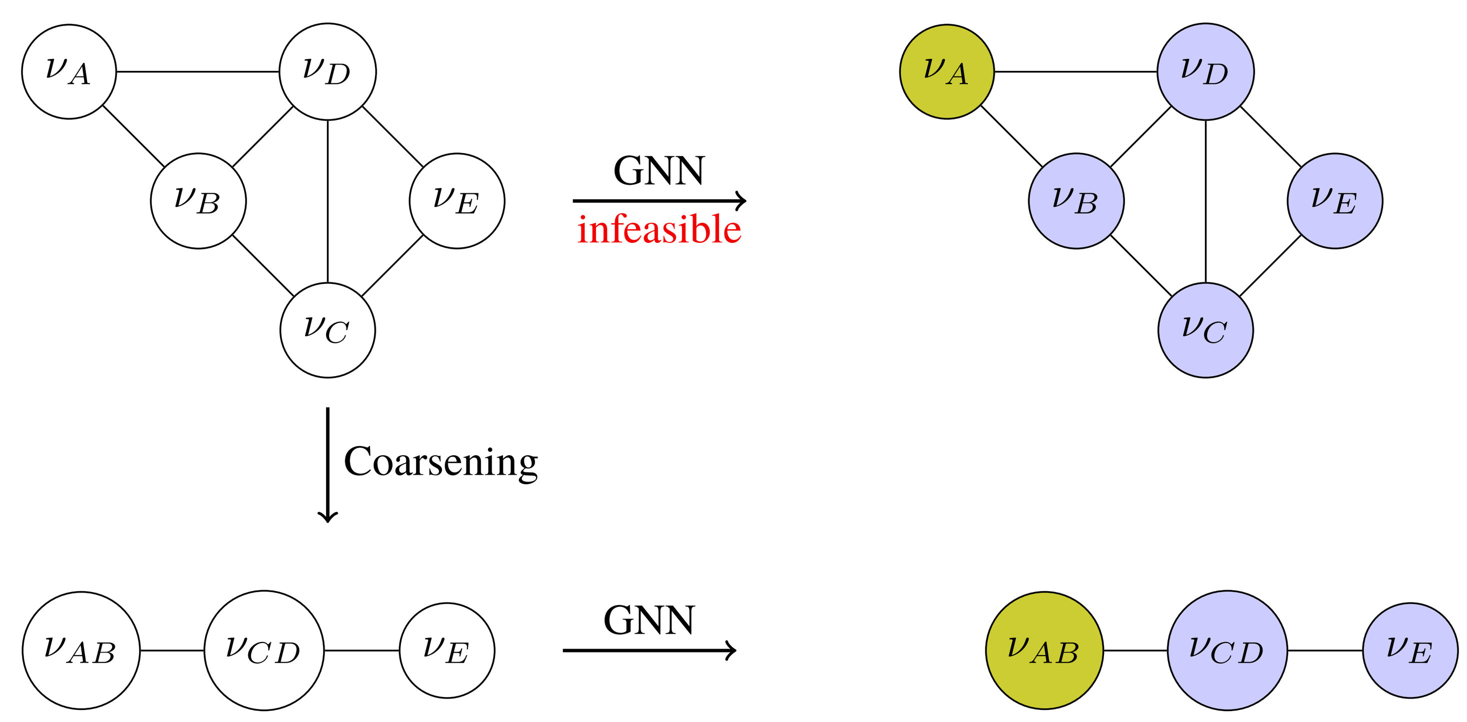

The performance-complexity trade-off

The main aim of our work is to explore the performance-complexity characteristics in the context of graph learning, as introduced in Procházka et al., 2022. The result of a repeated application of graph coarsening is a sequence of graphs $G_0, G_1, G_2, \dots, G_L$, where $G_0 = G$. Given a model that operates on graphs, a performance metric, and a complexity metric, the sequence $G_0, G_1, \dots, G_L$ corresponds to points in the performance-complexity plane, where advancing along the sequence generally hurts performance and decreases complexity. This performance-complexity characteristic allows for a choice of a working point that is optimal for the particular use-case.

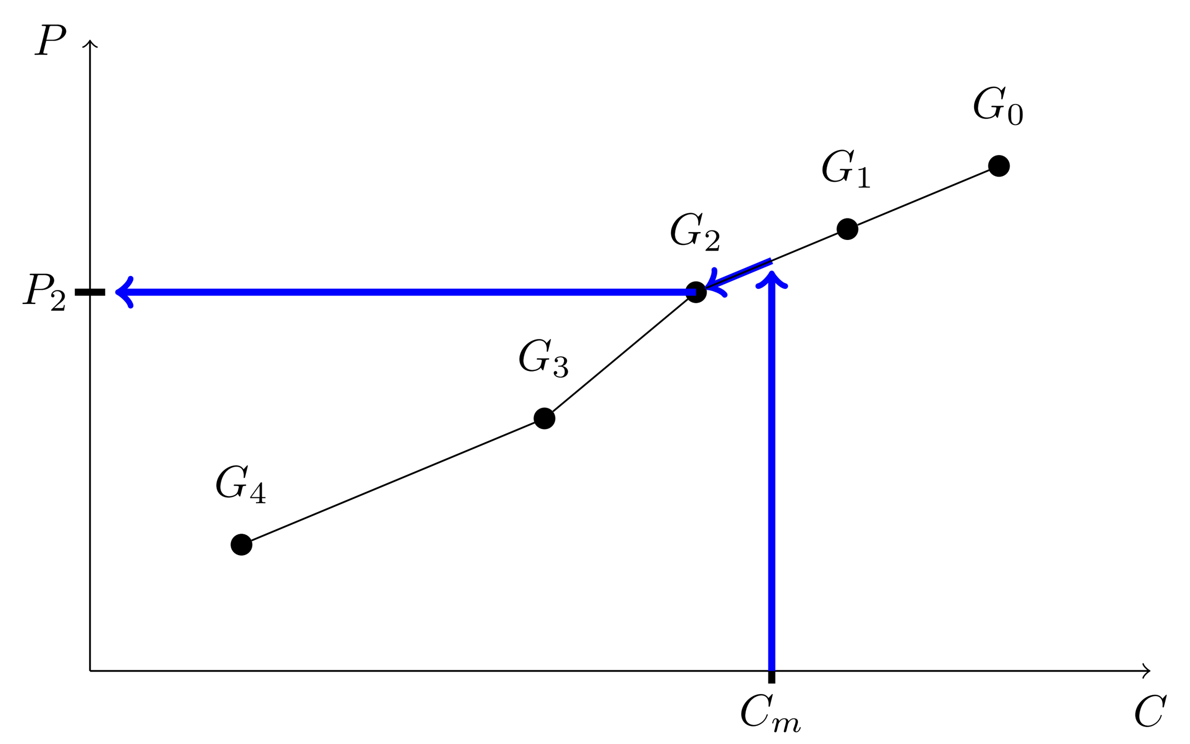

Achieving best possible performance P_2 for the complexity budget C_m.

HARP pipeline overview

Our work builds on the HARP method (see Chen et al., 2018) for pre-training on coarsened graphs. In HARP, the graph is coarsened using a pre-defined schema and the embedding is trained on the coarser graph first, prolonged to the original graph and finally trained on that as well.

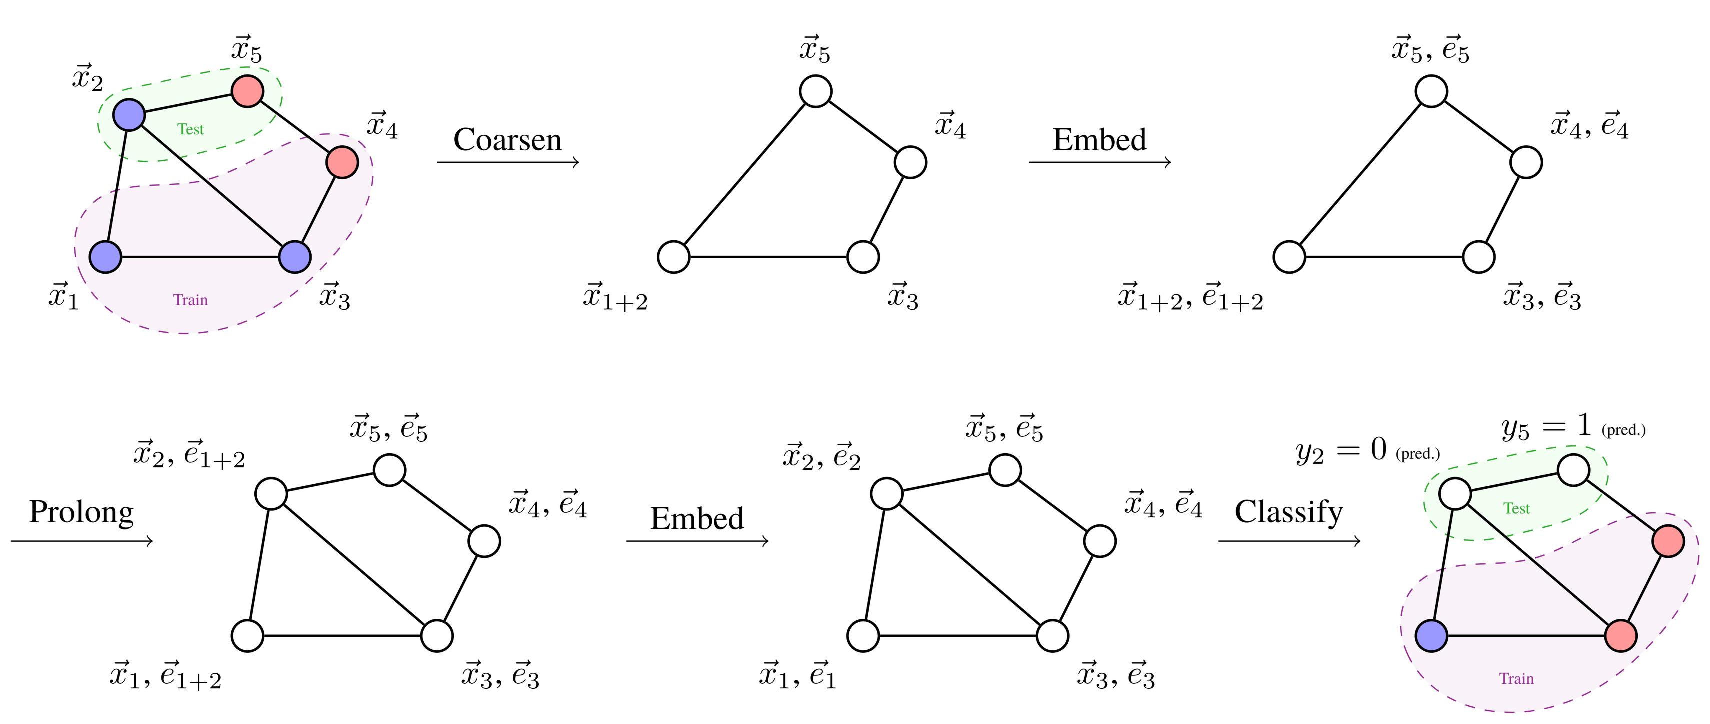

Deep HARP

The one-level HARP schema is generalized to obtain the sequence $G_0, G_1, G_2, \dots, G_L$ by repeatedly coarsening the graph. Subsequently, the method consists of alternating steps of training the embedding $\Phi_{G_{i+1}}$ on an intermediary graph G_{i+1}, followed by a prolongation to G_i, going from the coarsest graph G_L to the original G_0.

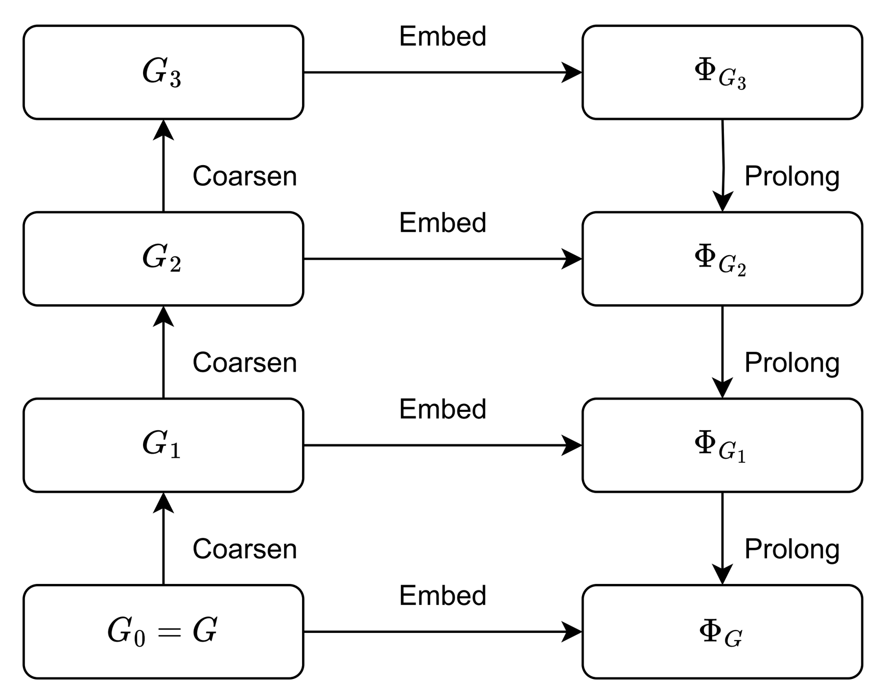

Adaptive prolongation

The main contribution of our work is the adaptive prolongation schema. The algorithm works with the pre-coarsened graphs produced by HARP, however, the prolongation steps are decoupled from the coarsening steps. The prolongation steps are driven by local properties of the graph and the downstream task, allowing for different levels of granularity in different parts of the graph. A single prolongation step is illustrated below.

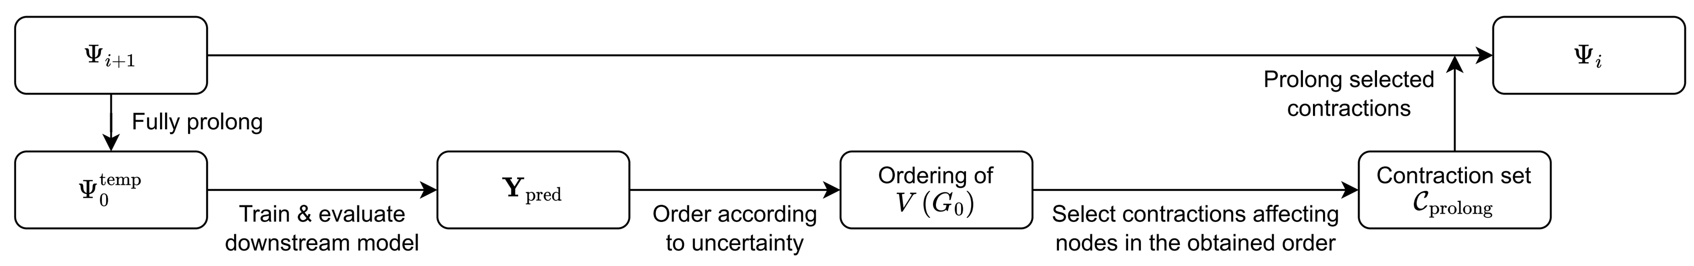

Results

Downstream classifier accuracies at different steps of adaptive prolongation. Dashed line shows the baseline node2vec model accuracy. The node count is taken relative to the total node count in each dataset. The results are averaged over multiple runs, with the solid line representing the mean and the shaded area denoting one standard deviation.

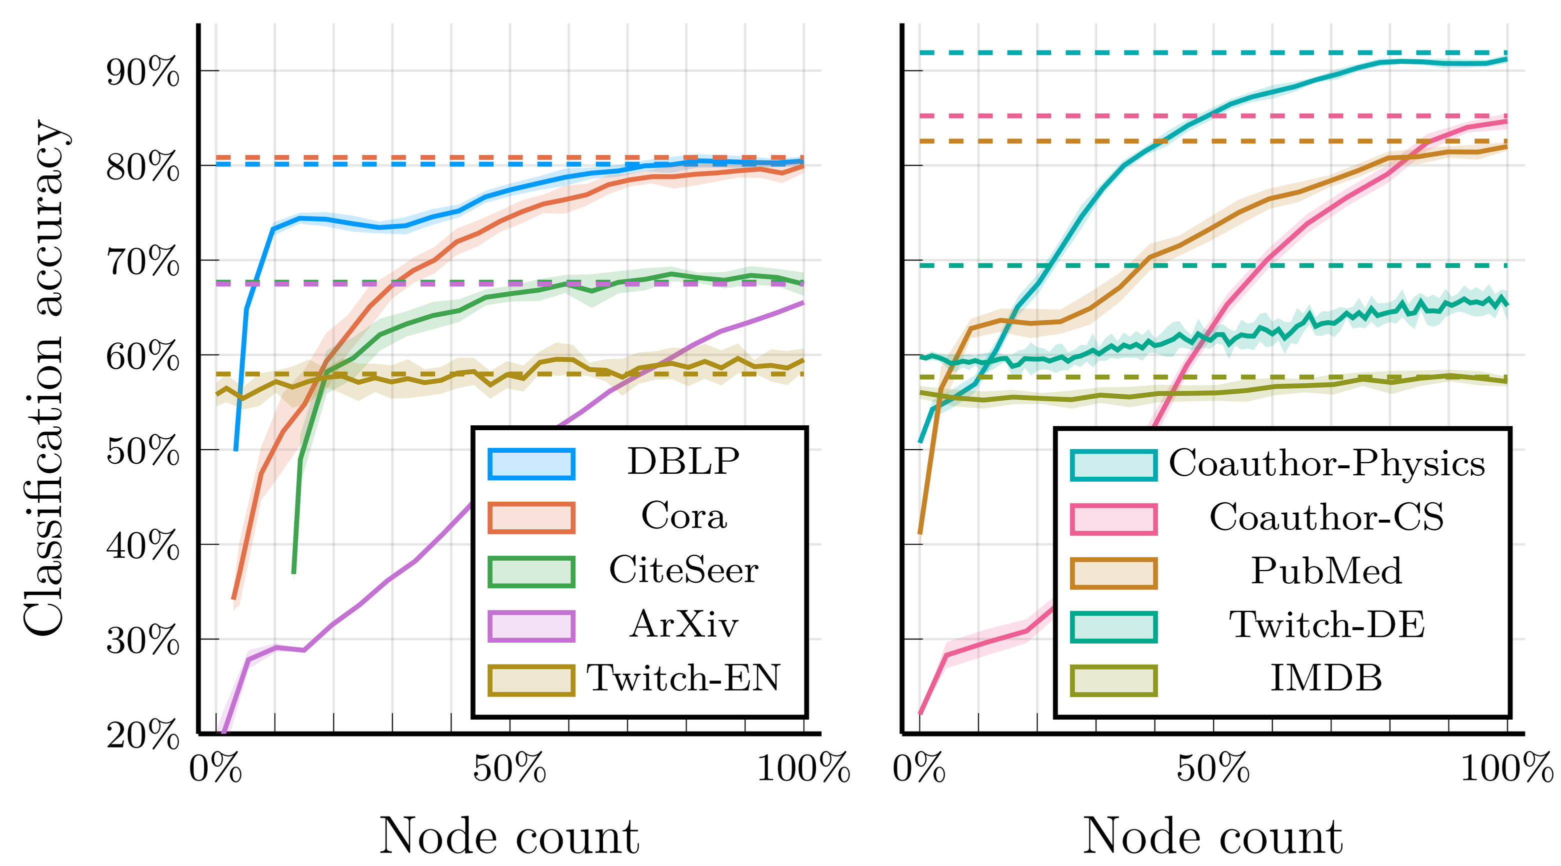

References

CHEN, Haochen; PEROZZI, Bryan; HU, Yifan; SKIENA, Steven, 2018. HARP: Hierarchical Representation Learning for Networks. *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 32, no. 1. ISSN 2374-3468.

PROCHÁZKA, Pavel; MAREŠ, Michal; DĚDIČ, Marek, 2022. Downstream Task Aware Scalable Graph Size Reduction for Efficient GNN Application on Big Data. In: *Information Technologies - Applications and Theory (ITAT 2022)*. Zuberec, Slovakia.